

# Randomized Dynamic Mode Decomposition

**N. Benjamin Erichson**  
University of St Andrews

**Steven L. Brunton**  
University of Washington

**J. Nathan Kutz**  
University of Washington

---

## Abstract

This paper presents a randomized algorithm for computing the near-optimal low-rank dynamic mode decomposition (DMD). Randomized algorithms are emerging techniques to compute low-rank matrix approximations. They are able to ease the computational challenges arising in the area of ‘big data’. The idea is to derive from the high-dimensional input matrix a smaller matrix, which is then used to efficiently compute the dynamic modes and eigenvalues. The algorithm is presented in a modular probabilistic framework, and the approximation quality can be controlled via oversampling, and power iterations.

*Keywords:* dynamic mode decomposition, randomized algorithm, dimension reduction.

---

## 1. Introduction

The dynamic mode decomposition is a dimensionality reduction technique, originally introduced in the field of fluid dynamics (Schmid 2010; Rowley, Mezić, Bagheri, Schlatter, and Henningson 2009). The method attempts to extract dynamic information from a dynamical system based on a sequence of snapshots (time series of data)  $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^m$ , separated in time by a constant step  $\Delta t$ . Specifically, the aim is to find the eigenvectors and eigenvalues of the time-independent linear map  $\mathbf{M} : \mathbb{R}^m \rightarrow \mathbb{R}^m$  which approximately maps a given snapshot  $\mathbf{x}_j$  onto the subsequent one  $\mathbf{x}_{j+1}$  as

$$\mathbf{x}_{j+1} \approx \mathbf{M}\mathbf{x}_j. \quad (1)$$

However, if the input data are high-dimensional, the linear map  $\mathbf{M}$  may be intractable to estimate and analyze directly. Instead, the singular value decomposition (SVD) is utilized to compute a rank-reduced representation of the linear map. This approach is robust to noise in the data and to numerical errors; however, the SVD is known to be computationally demanding. Hence, the computational time, and resources involved in computing the DMD can be tremendous for massive data matrices. The following two probabilistic strategies have been proposed previously to ease some of the computational challenges:

- (a) Erichson and Donovan (2016), and later Bistrian and Navon (2016) have presented an algorithm utilizing the randomized singular value decomposition (rSVD). While this approach is robust to noise, and reliable, only the computation of the SVD is accelerated. Subsequent computational steps involved in computing the DMD remain relatively

expensive. Hence, this approach is not entirely qualified to handle massive data.

- (b) Brunton, Proctor, Tu, and Kutz (2015), and later Erichson, Brunton, and Kutz (2016a) have presented the compressed dynamic mode decomposition, an algorithm which forms a smaller (sketched) sequence of snapshots from a small number of random linear combinations of the rows of the high-dimensional sequence of snapshots. Subsequently, the approximate dynamic modes and eigenvalues are obtained from the sketched representation. This approach achieves substantial speedups over (a), and has been shown to be successful for high-performance computing in the area of fluid flows, and video; however, the algorithm is less reliable, and is relatively sensitive to noise.

In the following, we propose a novel randomized algorithm for computing the near-optimal dynamic mode decomposition. Specifically, we are embedding the DMD into the probabilistic framework for computing low-rank matrix approximations as formulated by Halko, Martinsson, and Tropp (2011). The principal concept is depicted in Figure 1. Our algorithm achieves considerable speedups over (a) as well as it is more robust, and reliable than (b). Specifically, the approximation error can be controlled via oversampling and additional power iterations. This allows the user to choose the optimal trade-off between computational time and accuracy. Further, the algorithm is simple to implement and is embarrassingly parallel.

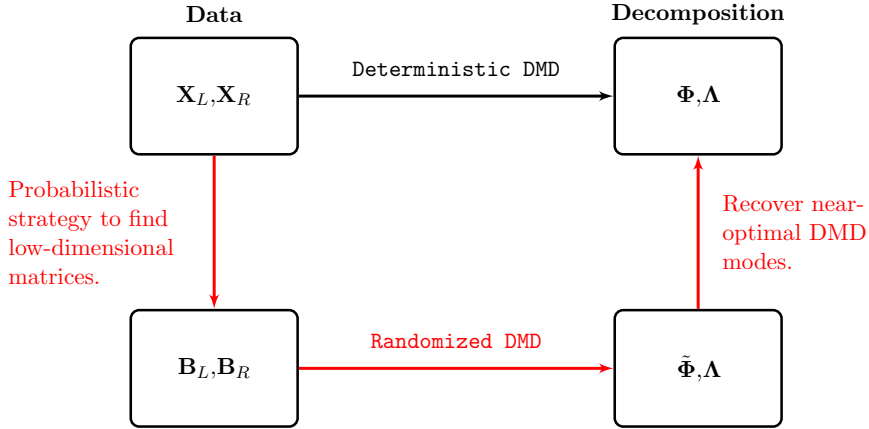


Figure 1: Conceptual architecture of the randomized dynamic mode decomposition (rDMD). First, smaller matrices are derived from the high-dimensional input data. The low-dimensional snapshot matrices  $\mathbf{B}_L$ , and  $\mathbf{B}_R$  are then used to compute the approximate dynamic modes  $\tilde{\Phi}$ , and eigenvalues  $\Lambda$ . Finally, the near-optimal modes may be recovered.

The rest of this manuscript is organized as follows. Section 2 briefly presents some notations and technical concepts used throughout the paper. Section 3 outlines the probabilistic framework utilized to compute the DMD. The randomized DMD algorithm is presented in Section 4. Section 5 shows some numerical results to demonstrate the performance of the proposed randomized DMD algorithm. Some final remarks and an outlook is given in Section 6.

## 2. Technical Preliminaries

In the following we present some notation as well as a brief review of the singular value decomposition (SVD) and the dynamic mode decomposition (DMD).

### 2.1. Notation

Vectors in  $\mathbb{R}^n$  and  $\mathbb{C}^n$  are denoted as bold lower case letters  $\mathbf{x} = [x_1, x_2, \dots, x_n]$ . Both real  $\mathbb{R}^{m \times n}$  and complex  $\mathbb{C}^{m \times n}$  matrices are denoted by bold capitals  $\mathbf{X}$ , and its entry at row  $i^{th}$  and column  $j^{th}$  is denoted as  $\mathbf{X}(i, j)$ . The transpose of a matrix is denoted as  $\mathbf{X}^*$ , and in case of a complex matrix this denotes the Hermitian transpose. The spectral or operator norm of a matrix is defined as the largest singular value  $\sigma_{max}(\mathbf{X})$  of the matrix  $\mathbf{X}$ , i.e., the square root of the largest eigenvalue  $\lambda_{max}(\mathbf{X})$  of the positive-semidefinite matrix  $\mathbf{X}^* \mathbf{X}$

$$\|\mathbf{X}\|_2 = \sqrt{\lambda_{max}(\mathbf{X}^* \mathbf{X})} = \sigma_{max}(\mathbf{X}).$$

The Frobenius norm is defined as the square root of the sum of the absolute squares of its elements, which is equal to the square root of the matrix trace of  $\mathbf{X}^* \mathbf{X}$

$$\|\mathbf{X}\|_F = \sqrt{\sum_{i=1}^m \sum_{j=1}^n |\mathbf{X}(i, j)|^2} = \sqrt{\text{trace}(\mathbf{X}^* \mathbf{X})}.$$

The relative reconstruction error is computed as  $\|\mathbf{X} - \hat{\mathbf{X}}\|_F / \|\mathbf{X}\|_F$ , where  $\hat{\mathbf{X}}$  denotes an approximated matrix. The column space (range) of  $\mathbf{X}$  is denoted as  $\text{col}(\mathbf{X})$ , and the row space as  $\text{row}(\mathbf{X})$ .

### 2.2. The Singular Value Decomposition

Given an  $m \times n$  matrix  $\mathbf{X}$ , the ‘economic’ singular value decomposition (SVD) admits the factorization

$$\mathbf{X} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*, \tag{2}$$

where the matrices  $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_r] \in \mathbb{R}^{m \times r}$  and  $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_n] \in \mathbb{R}^{r \times n}$  are orthonormal,  $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$  is a diagonal, and  $r = \min(m, n)$ . The  $r$  left singular vectors in  $\mathbf{U}$  provide a basis for the column space, and the  $r$  right singular vectors in  $\mathbf{V}$  a basis for the domain of the matrix  $\mathbf{X}$ .  $\mathbf{\Sigma}$  contains the corresponding non-negative singular values  $\sigma_1 \geq \dots \geq \sigma_r \geq 0$ , describing the spectrum of the data. Often, only the  $k$  dominant singular vectors and values are of interest. The low-rank SVD of rank  $k$  takes the form:

$$\mathbf{X}_k = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^* = [\mathbf{u}_1, \dots, \mathbf{u}_k] \text{diag}(\sigma_1, \dots, \sigma_k) [\mathbf{v}_1, \dots, \mathbf{v}_k]^*.$$

**The Moore-Penrose Pseudoinverse.** Given the singular value decomposition  $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*$ , the pseudoinverse  $\mathbf{X}^\dagger$  is computed as

$$\mathbf{X}^\dagger = \mathbf{V}\mathbf{\Sigma}^{-1}\mathbf{U}^*. \quad (3)$$

We use the Moore-Penrose pseudoinverse in the following to provide a least squares solution to a system of linear equations.

### 2.3. Dynamic Mode Decomposition

The objective of the dynamic mode decomposition is to find an estimate for the linear map  $\mathbf{M}$  and its eigenvalue decomposition. Following [Tu, Rowley, Luchtenberg, Brunton, and Kutz \(2014\)](#), the deterministic DMD algorithm proceeds by first separating the snapshot sequence  $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^m$  into two overlapping sets of data

$$\mathbf{X}_L = \begin{bmatrix} | & | & & | \\ \mathbf{x}_0 & \mathbf{x}_1 & \cdots & \mathbf{x}_{n-1} \\ | & | & & | \end{bmatrix}, \quad \mathbf{X}_R = \begin{bmatrix} | & | & & | \\ \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_n \\ | & | & & | \end{bmatrix}. \quad (4)$$

$\mathbf{X}_L \in \mathbb{R}^{m \times n}$  and  $\mathbf{X}_R \in \mathbb{R}^{m \times n}$  are called the left and right snapshot sequences. Next, Equation (1) can be reformulated in matrix notation

$$\mathbf{X}_R \approx \mathbf{M}\mathbf{X}_L. \quad (5)$$

Then, in order to find an estimate for the linear map  $\mathbf{M}$ , the following least-squares problem can be formulated

$$\hat{\mathbf{M}} = \underset{\mathbf{M}}{\operatorname{argmin}} \|\mathbf{X}_R - \mathbf{M}\mathbf{X}_L\|_F^2. \quad (6)$$

Specifically, the low-rank DMD algorithm proceeds as follows:

1. Given a target rank  $k$ , compute the truncated singular value decomposition of the left snapshot matrix  $\mathbf{X}_L$ :

$$\mathbf{X}_L = \mathbf{U}_k \mathbf{\Sigma}_k \mathbf{V}_k^*, \quad (7)$$

where  $\mathbf{U}_k \in \mathbb{R}^{m \times k}$ , and  $\mathbf{V}_k \in \mathbb{R}^{n \times k}$  are the truncated left and right singular vectors.  $\mathbf{\Sigma}_k \in \mathbb{R}^{k \times k}$  contains the corresponding singular values.

2. The estimator for the best-fit linear map is defined as

$$\hat{\mathbf{M}} := \mathbf{X}_R \mathbf{X}_L^\dagger = \mathbf{X}_R \mathbf{V}_k \mathbf{\Sigma}_k^{-1} \mathbf{U}_k^*. \quad (8)$$

However, if the data are high-dimensional (i.e.  $m$  is large) this estimator tends to be infeasible to compute. Instead, the linear map is projected onto the dominant  $k$  left

singular vectors as follows

$$\tilde{\mathbf{M}} = \mathbf{U}_k^* \hat{\mathbf{M}} \mathbf{U}_k = \mathbf{U}_k^* \mathbf{X}_R \mathbf{V}_k \boldsymbol{\Sigma}_k^{-1}. \quad (9)$$

This yields the much smaller matrix  $\tilde{\mathbf{M}} \in \mathbb{R}^{k \times k}$ . Specifically,  $\tilde{\mathbf{M}}$  and  $\hat{\mathbf{M}}$  are related via the similarity transform. Hence,  $\tilde{\mathbf{M}}$  features the same dominant eigenvalues as  $\hat{\mathbf{M}}$  does. Further, it is simple to show that the dominant eigenvectors of  $\hat{\mathbf{M}}$  can be recovered from  $\tilde{\mathbf{M}}$  as well.

3. Next, the eigendecomposition of  $\tilde{\mathbf{M}}$  is computed

$$\tilde{\mathbf{M}} \mathbf{W} = \mathbf{W} \boldsymbol{\Lambda}, \quad (10)$$

where columns of  $\mathbf{W} \in \mathbb{C}^{k \times k}$  are eigenvectors, and  $\boldsymbol{\Lambda} \in \mathbb{C}^{k \times k}$  is a diagonal matrix containing the corresponding complex eigenvalues  $\lambda_j$ . The continuous-time eigenvalues are given by  $\omega_j = \log(\lambda_j)/\Delta t$ .

4. Finally, we may recover the dominant  $k$  DMD modes (eigenvectors) of  $\hat{\mathbf{M}}$  as follows

$$\boldsymbol{\Phi}_k = \mathbf{X}_R \mathbf{V}_k \boldsymbol{\Sigma}_k^{-1} \mathbf{W}_k, \quad (11)$$

where  $\boldsymbol{\Phi} \in \mathbb{C}^{m \times k}$ . Note that Equation (11) from Tu *et al.* (2014) differs from the formula  $\boldsymbol{\Phi} = \mathbf{U} \mathbf{W}$  from Schmid (2010), although these will tend to converge if  $\mathbf{X}_L$  and  $\mathbf{X}_R$  have the same column spaces.

### 3. Probabilistic Framework

Over the past two decades, probabilistic algorithms to compute low-rank matrix approximations have become prominent (Halko *et al.* 2011; Mahoney 2011; Drineas and Mahoney 2016; Martinsson 2016). Randomness is used as a computational strategy to find a smaller representation, derived from a high-dimensional input data matrix  $\mathbf{X} \in \mathbb{R}^{m \times n}$ . This smaller matrix can then be used to compute a low-rank approximation. Following Halko *et al.* (2011), this probabilistic framework can be split into two computational stages:

- **Stage A:** Find a near-optimal orthonormal basis  $\mathbf{Q} \in \mathbb{R}^{m \times k}$  for the range of the input matrix  $\mathbf{X}$ , such that  $\mathbf{X} \approx \mathbf{Q} \mathbf{Q}^* \mathbf{X}$  is satisfied.
- **Stage B:** Given the near-optimal basis  $\mathbf{Q}$ , restrict the input matrix to a low-dimensional space  $\mathbf{B} \in \mathbb{R}^{k \times n}$ . This smaller matrix can then be used to compute the near-optimal low-rank dynamic mode decomposition.

### 3.1. Stage A: Computing a Near-Optimal Basis

The first stage is involved in approximating the range of the input matrix. Given a target rank  $k \ll (m, n)$ , the aim is to compute a near-optimal basis  $\mathbf{Q} \in \mathbb{R}^{m \times k}$  for the input matrix  $\mathbf{X} \in \mathbb{R}^{m \times n}$  such that

$$\mathbf{X} \approx \mathbf{Q}\mathbf{Q}^*\mathbf{X} \quad (12)$$

is satisfied. Specifically, the range of the high-dimensional input matrix is sampled using the concept of random projections. Thus a basis is efficiently computed as

$$\mathbf{Y} = \mathbf{X}\mathbf{\Omega}, \quad (13)$$

where  $\mathbf{\Omega} \in \mathbb{R}^{n \times k}$  denotes a random test matrix which is drawn from a sub-Gaussian distribution. An orthonormal basis  $\mathbf{Q} \in \mathbb{R}^{m \times k}$  is then obtained via the QR-decomposition  $\mathbf{Y} = \mathbf{Q}\mathbf{R}$ , such that Equation (12) is satisfied.

### 3.2. Stage B: Computing a Low-Dimensional Matrix

Next, we derive a smaller matrix from the high-dimensional input matrix. Specifically, given the near-optimal basis  $\mathbf{Q}$ , the matrix  $\mathbf{X}$  is projected to a low-dimensional space

$$\mathbf{B} = \mathbf{Q}^*\mathbf{X}, \quad (14)$$

so that we yield the smaller matrix  $\mathbf{B} \in \mathbb{R}^{k \times n}$ . This process preserves the geometric structure in an Euclidean sense, i.e., the angles between vectors and their length are preserved. It follows that

$$\mathbf{X} \approx \mathbf{Q}\mathbf{B}. \quad (15)$$

Subsequently, the low-dimensional matrix  $\mathbf{B}$  can be used to construct a low-rank approximation.

### 3.3. Computational Considerations

The computational steps involved in computing the approximate basis  $\mathbf{Q}$  and the low-dimensional matrix  $\mathbf{B}$  are simple to implement, and embarrassingly parallel. Thus, randomized algorithms can benefit from modern computational architectures, and they are in particular suitable for GPU accelerated computing. A favorable computational aspect is also that only two passes over the input matrix are required in order to obtain the low-dimensional matrix. Pass-efficiency is a crucial aspect when dealing with massive data matrices, which are too big to fit into fast memory. Reading data from hard disk is prohibitively slow, and often poses the actual bottleneck. Further, the performance of the basis  $\mathbf{Q}$  can be improved using the concept of oversampling and the power scheme as outlined in the following.

### Oversampling

In theory, if the matrix  $\mathbf{X}$  has exact rank  $k$ , the sampled matrix  $\mathbf{Y}$  spans with high probability a basis for the column space. In practice, however, it is common that the singular values  $\{\sigma_i\}_{i=k+1}^n$  are non-zero. Thus, it is favorable to construct a slightly bigger test matrix in order to obtain an improved basis. This means, we compute  $\mathbf{Y}$  using an  $\mathbf{\Omega} \in \mathbb{R}^{n \times l}$  test matrix instead, where  $l = k + p$ . Thus, the oversampling parameter  $p$  denotes the number of additional samples. In most situations small values  $p = \{5, 10\}$  are sufficient to obtain a good basis that is comparable to the best possible basis (Halko *et al.* 2011; Martinsson 2016).

### Power Scheme

A second strategy in order to improve the performance of the basis is the concept of power iterations (Rokhlin, Szlam, and Tygert 2009; Martinsson 2016; Gu 2015). In particular, a slowly decaying singular value spectrum of the input matrix can seriously affect the quality of the approximated basis matrix  $\mathbf{Q}$ . Thus, the method of power iterations is used to pre-process the input matrix in order to enforce a faster decaying singular value spectrum. Specifically, the sampling matrix  $\mathbf{Y}$  is obtained as follows

$$\mathbf{Y} = ((\mathbf{X}\mathbf{X}^*)^q \mathbf{X})\mathbf{\Omega} \quad (16)$$

where  $q$  is an integer specifying the number of power iterations. It can be showed that  $\mathbf{X}^{(q)} := (\mathbf{X}\mathbf{X}^*)^q \mathbf{X} = \mathbf{U}\mathbf{\Sigma}^{2q+1}\mathbf{V}^*$ . Hence, for  $q > 0$ , the pre-processed matrix  $\mathbf{X}^{(q)}$  has a relatively fast decay of singular values compared to the input matrix  $\mathbf{X}$ . The drawback of this method is that additional passes over the input matrix are required. However, as few as  $q = \{1, 2\}$  power iterations can considerably improve the approximation quality, even when the singular values of the input matrix decay slowly.

### Theoretical Performance

Both the concept of oversampling and the power scheme allow to control the quality of low-rank approximation. Martinsson (2016) provides the following description of the average case behavior of the outlined probabilistic framework: <sup>1</sup>

$$\mathbb{E}\|\mathbf{X} - \mathbf{Q}\mathbf{Q}^*\mathbf{X}\|_2 \leq \left[ 1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{l}}{p} \cdot \sqrt{\min\{m, n\} - k} \right]^{\frac{1}{2q+1}} \sigma_{k+1}(\mathbf{X}).$$

Here it is assumed that  $p \geq 2$ . Thus, both oversampling and the computation of additional power iterations drive the approximation error down.

---

<sup>1</sup>Note, that this is a simplified version of one of the key theorems presented by Halko *et al.* (2011), who provide a detailed error analysis of the outlined probabilistic framework.

#### 4. Randomized Dynamic Mode Decomposition

In the following we present a novel algorithm to compute the dynamic mode decomposition. The algorithm as presented is simple to implement, and can fully benefit from modern computational architectures, e.g., parallelization and multi-threading.

Given a sequence of snapshots  $\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^m$ , we first compute the near-optimal basis  $\mathbf{Q} \in \mathbb{R}^{m \times l}$ . Then, we project the data onto the low-dimensional space, so that we obtain the low-dimensional sequence of snapshots  $\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_n \in \mathbb{R}^l$ . Here,  $l = k + p$ , where  $k$  denotes the desired target rank, and  $p$  the oversampling parameter. Next, we separate the sequence into two overlapping matrices  $\mathbf{B}_L \in \mathbb{R}^{l \times n}$  and  $\mathbf{B}_R \in \mathbb{R}^{l \times n}$  similar to Equation (4) as

$$\mathbf{B}_L = \begin{bmatrix} | & | & & | \\ \mathbf{b}_0 & \mathbf{b}_1 & \cdots & \mathbf{b}_{n-1} \\ | & | & & | \end{bmatrix}, \quad \mathbf{B}_R = \begin{bmatrix} | & | & & | \\ \mathbf{b}_1 & \mathbf{b}_2 & \cdots & \mathbf{b}_n \\ | & | & & | \end{bmatrix}. \quad (17)$$

Now, the least-squares problem in Equation (6) can be re-formulated as

$$\hat{\mathbf{M}}_{\mathbf{B}} = \underset{\mathbf{M}_{\mathbf{B}}}{\operatorname{argmin}} \|\mathbf{B}_R - \mathbf{M}_{\mathbf{B}} \mathbf{B}_L\|_F^2, \quad (18)$$

The following computational steps are then equivalent to the deterministic DMD algorithm. Using the Moore-Penrose pseudoinverse, the estimator for the linear map  $\hat{\mathbf{M}}_{\mathbf{B}} \in \mathbb{R}^{l \times l}$  is defined as

$$\hat{\mathbf{M}}_{\mathbf{B}} := \mathbf{B}_R \mathbf{B}_L^\dagger = \mathbf{B}_R \mathbf{V} \mathbf{\Sigma}^{-1} \tilde{\mathbf{U}}^*, \quad (19)$$

where the matrices  $\mathbf{U} \in \mathbb{R}^{l \times k}$ , and  $\mathbf{V} \in \mathbb{R}^{n \times k}$  are the truncated left and right singular vectors. The diagonal matrix  $\mathbf{\Sigma} \in \mathbb{R}^{k \times k}$  has the corresponding singular values as entries. Note, that if  $p = 0$ , then  $l = k$ , so that truncation is redundant.<sup>2</sup> Next,  $\hat{\mathbf{M}}_{\mathbf{B}}$  is projected onto the left singular vectors

$$\tilde{\mathbf{M}}_{\mathbf{B}} = \tilde{\mathbf{U}}^* \hat{\mathbf{M}}_{\mathbf{B}} \tilde{\mathbf{U}} \quad (20a)$$

$$= \tilde{\mathbf{U}}^* \mathbf{B}_R \mathbf{V} \mathbf{\Sigma}^{-1}. \quad (20b)$$

The DMD modes, containing the spatial information, are then obtained by computing the eigendecomposition of  $\tilde{\mathbf{M}}_{\mathbf{B}} \in \mathbb{R}^{k \times k}$

$$\tilde{\mathbf{M}}_{\mathbf{B}} \tilde{\mathbf{W}}_{\mathbf{B}} = \tilde{\mathbf{W}}_{\mathbf{B}} \mathbf{\Lambda}, \quad (21)$$

where the columns of  $\tilde{\mathbf{W}}_{\mathbf{B}} \in \mathbb{C}^{k \times k}$  are eigenvectors  $\phi_j$ , and  $\mathbf{\Lambda} \in \mathbb{C}^{k \times k}$  is a diagonal matrix

---

<sup>2</sup>Alternatively, the computations in the following can be performed without truncating the SVD. This might increase the performance slightly, but it comes with higher computational costs. Then the dynamic modes are truncated at the end.



containing the corresponding eigenvalues  $\lambda_j$ . The randomized DMD modes  $\Phi_B \in \mathbb{C}^{l \times k}$  are consequently given by

$$\Phi_B = B_R V \Sigma^{-1} \tilde{W}_B. \quad (22)$$

Finally, the high-dimensional DMD modes  $\Phi \in \mathbb{C}^{m \times k}$  may be recovered as

$$\Phi = Q \Phi_B = Q B_R V \Sigma^{-1} \tilde{W}_B, \quad (23)$$

Note that Equation (23) follows Tu *et al.* (2014), and differs from the formula  $\Phi = Q \tilde{U} \tilde{W}_B$  which follows Schmid (2010). The computational steps for a practical implementation are sketched in Algorithm 1. An implementation in Python is available via the GIT repository <https://github.com/Benli11/DMDpack>.

---

**Algorithm 1** Randomized Dynamic Mode Decomposition (rDMD).

Given a snapshot matrix  $X \in \mathbb{R}^{m \times n}$ , and a target rank  $k \leq \min(m, n)$ , the near-optimal dominant dynamic modes  $\Phi \in \mathbb{C}^{m \times k}$  and eigenvalues  $\Lambda \in \mathbb{C}^{k \times k}$  are computed. The approximation quality can be controlled via oversampling  $p$  and the computation of power iterations  $q$ .

---

- |      |  |   |
|------|--|---|
| (1)  | $l = k + p$                                      | Slight oversampling.                        |
| (2)  | $\Omega = \text{rand}(n, l)$                     | Generate random test matrix.                |
| (3)  | $Y = X\Omega$                                    | Compute sampling matrix.                    |
| (4)  | <b>for</b> $j = 1, \dots, q$                     | Power iterations (optional).                |
| (5)  | $[Q, \sim] = \text{qr}(Y)$                       |   |
| (6)  | $[Z, \sim] = \text{qr}(X^*Q)$                    |   |
| (7)  | $Y = XZ$   |   |
| (8)  | <b>end for</b>                                   |   |
| (9)  | $[Q, \sim] = \text{qr}(Y)$                       | Orthonormalize sampling matrix.             |
| (10) | $B = Q^*X$                                       | Project input matrix to smaller space.      |
| (11) | $B_L, B_R = B$                                   | Left/right low-dimensional snapshot matrix. |
| (12) | $[U, \Sigma, V] = \text{svd}(B_L, k)$            | Truncated SVD.                              |
| (13) | $\tilde{M}_B = U^* B_R V \Sigma^{-1}$            | Least squares fit.                          |
| (14) | $[\tilde{W}, \Lambda] = \text{eig}(\tilde{M}_B)$ | Eigenvalue decomposition.                   |
| (15) | $\Phi \leftarrow Q B_R V \Sigma^{-1} \tilde{W}$  | Recover high-dimensional DMD modes $\Phi$ . |
- 

*Remark 1.* A standard approach to construct the test matrix  $\Omega$  is to draw random samples from the Gaussian or uniform distribution. However, the multiplication of dense matrices can become computational expensive. An alternative is to construct a structured random test matrix which allows fast matrix multiplications (Woolfe, Liberty, Rokhlin, and Tygert 2008).

*Remark 2.* A direct implementation of the power scheme, as outlined in Section 3.3.2, is numerically unstable due to round-off errors. Instead, the sampling matrix  $Y$  is orthogonalized in between each computational step to improve the stability (Martinsson 2016).

*Remark 3.* As default values for the oversampling and power iteration parameter, we suggest  $p = 10$ , and  $q = 1$ , see also Erichson, Voronin, Brunton, and Kutz (2016b).

*Justification*

In the following we outline some justification for the proposed algorithm. Specifically, we want to show that the dominant  $k$  eigenvectors and eigenvalues of  $\hat{\mathbf{M}}$  and  $\hat{\mathbf{M}}_{\mathbf{B}}$  can be approximately related as

$$\hat{\mathbf{M}} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^* \approx \mathbf{Q}\mathbf{W}_{\mathbf{B}}\mathbf{\Lambda}_B\mathbf{W}_{\mathbf{B}}^*\mathbf{Q}^* = \mathbf{Q}\hat{\mathbf{M}}_{\mathbf{B}}\mathbf{Q}^*, \quad (24)$$

given an orthonormal basis  $\mathbf{Q}$  so that

$$\mathbf{X}_L \approx \mathbf{Q}\mathbf{Q}^*\mathbf{X}_L \quad (25a)$$

$$\mathbf{X}_R \approx \mathbf{Q}\mathbf{Q}^*\mathbf{X}_R, \quad (25b)$$

is satisfied. Starting with Equation (8), we have defined the estimator for  $\hat{\mathbf{M}}$  as

$$\mathbf{W}\mathbf{\Lambda}\mathbf{W}^* = \hat{\mathbf{M}} := \mathbf{X}_R\mathbf{X}_L^\dagger = \mathbf{X}_R\mathbf{V}_k\Sigma_k^{-1}\mathbf{U}_k^*.$$

Now, substituting Equation (25a) and (25b) yields the following approximate relationship

$$\mathbf{W}\mathbf{\Lambda}\mathbf{W}^* = \hat{\mathbf{M}} \approx (\mathbf{Q}\mathbf{Q}^*\mathbf{X}_R)(\mathbf{Q}\mathbf{Q}^*\mathbf{X}_L)^\dagger = (\mathbf{Q}\mathbf{Q}^*\mathbf{X}_R)(\mathbf{V}_k\Sigma_k^{-1}\mathbf{U}_k^*\mathbf{Q}\mathbf{Q}^*). \quad (26)$$

This can be rewritten by substituting  $\mathbf{B}_L := \mathbf{Q}^*\mathbf{X}_L$  and  $\mathbf{B}_R := \mathbf{Q}^*\mathbf{X}_R$

$$\mathbf{W}\mathbf{\Lambda}\mathbf{W}^* = \hat{\mathbf{M}} \approx (\mathbf{Q}\mathbf{B}_R)(\mathbf{Q}\mathbf{B}_L)^\dagger = (\mathbf{Q}\mathbf{B}_R)(\mathbf{V}_k\Sigma_k^{-1}\tilde{\mathbf{U}}_k^*\mathbf{Q}^*). \quad (27)$$

Next, we note that Equation (19) defines  $\hat{\mathbf{M}}_{\mathbf{B}} := \mathbf{B}_R\mathbf{V}\Sigma^{-1}\tilde{\mathbf{U}}^*$ . Then, substitution yields

$$\mathbf{W}\mathbf{\Lambda}\mathbf{W}^* = \hat{\mathbf{M}} \approx \mathbf{Q}\hat{\mathbf{M}}_{\mathbf{B}}\mathbf{Q}^* = \mathbf{Q}\mathbf{W}_{\mathbf{B}}\mathbf{\Lambda}_B\mathbf{W}_{\mathbf{B}}^*\mathbf{Q}^*. \quad (28)$$

Hence, we have  $\mathbf{W} \approx \mathbf{Q}\mathbf{W}_{\mathbf{B}}$  and  $\mathbf{\Lambda} \approx \mathbf{\Lambda}_B$ . Note, that the approximation depends on the quality of the orthonormal basis  $\mathbf{Q}$ .

It is then easy to verify that the same results approximately hold for  $\tilde{\mathbf{M}}_{\mathbf{B}}$ , i.e., the linear map  $\hat{\mathbf{M}}_{\mathbf{B}}$  projected onto the left singular vectors. Specifically, we have

$$\tilde{\mathbf{U}}^*\hat{\mathbf{M}}_{\mathbf{B}}\tilde{\mathbf{U}} = \tilde{\mathbf{U}}^*\mathbf{W}_{\mathbf{B}}\mathbf{\Lambda}_B\mathbf{W}_{\mathbf{B}}^*\tilde{\mathbf{U}}. \quad (29)$$

Then  $\tilde{\mathbf{W}}_{\mathbf{B}} := \tilde{\mathbf{U}}^*\mathbf{W}_{\mathbf{B}}$ , and similar to the previous reasoning it follows that

$$\hat{\mathbf{M}} = \mathbf{W}\mathbf{\Lambda}\mathbf{W}^* \approx \mathbf{Q}\tilde{\mathbf{U}}\tilde{\mathbf{W}}_{\mathbf{B}}\mathbf{\Lambda}_B\tilde{\mathbf{W}}_{\mathbf{B}}^*\tilde{\mathbf{U}}^*\mathbf{Q}^*. \quad (30)$$

Hence, we have  $\mathbf{W} \approx \mathbf{Q}\tilde{\mathbf{U}}\tilde{\mathbf{W}}_{\mathbf{B}}$  and  $\mathbf{\Lambda} \approx \mathbf{\Lambda}_B$ .

## 5. Numerical Results

In the following we present numerical results demonstrating the computational performance of the randomized dynamic mode decomposition (rDMD). All computations are performed on a standard notebook (Intel Core i7-5500U 2.4GHz, 8GB DDR3 memory). The underlying numerical linear algebra routines are accelerated using the Intel Math Kernel Library (MKL).

As a canonical example we use a fluid flow behind a cylinder (Noack, Afanasiev, Morzynski, Tadmor, and Thiele 2003). Specifically, the data are constructed as a sequence of 151 snapshots of fluid vorticity fields behind a stationary cylinder on an equispaced  $449 \times 199$  grid.<sup>3</sup> The flow features a periodically shedding wake structure at Reynolds number  $Re = 100$ , and is inherently low-rank. In the following we are interested in computing the first  $k = 15$  dominant DMD modes and eigenvalues. In order to compute the randomized DMD we chose an oversampling parameter  $p = 10$ , which we suggest as the default value for oversampling. Knowing that the data are inherently low-rank no additional power iterations are computed. However, we suggest as default parameter  $q = 1$ . Figure 2 shows the first six dominant DMD modes computed using both the deterministic and the randomized DMD algorithm.

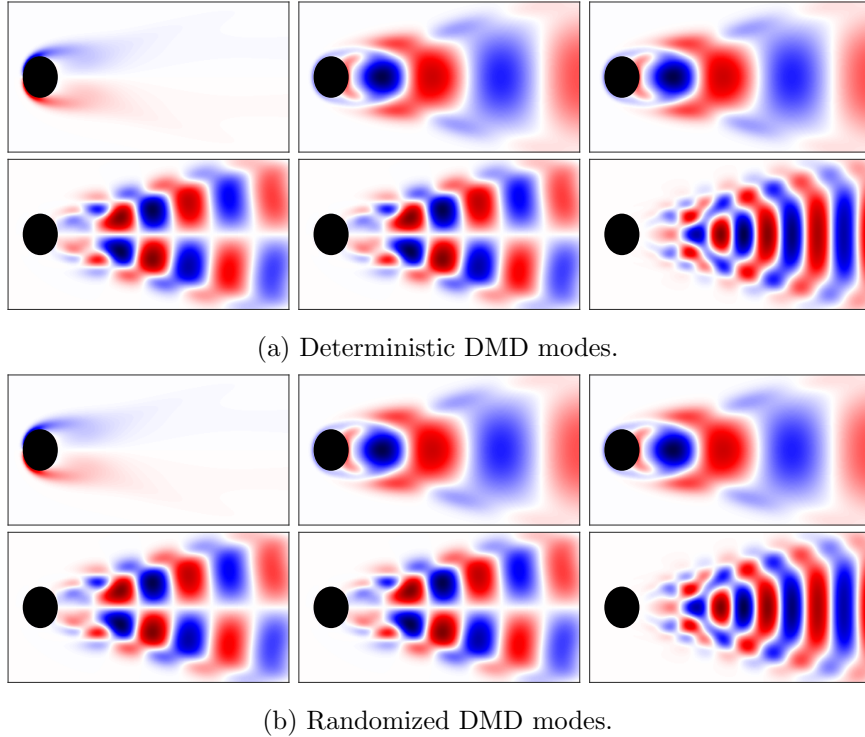


Figure 2: Dominant six dynamic modes of a fluid flow behind a cylinder.

<sup>3</sup>The data are obtained by solving the two-dimensional Navier-Stokes equations using the publicly available code at <https://github.com/cwrowley/ibpm>, based on the immersed boundary projection method in a fast multi-domain solver (Taira and Colonius 2007; Colonius and Taira 2008).

Note, that the data are not mean centered. Hence, the first mode does not change over time, and corresponds to the zero frequency DMD eigenvalue. The randomized algorithm faithfully reveals the coherent structures, while requiring considerably less computational resources. Specifically, a 6 fold speedup is achieved over the deterministic algorithm. The reconstruction error of the deterministic algorithm is  $5.11e - 03$ , while the randomized algorithm achieves an error as low as  $5.17e - 03$ .

Figure 3a shows the dominant eigenvalues. The proposed randomized DMD algorithm, as well as the compressed DMD (cDMD) algorithm (Brunton *et al.* 2015) and the DMD algorithm utilizing the randomized singular value decomposition (Erichson and Donovan 2016; Bistrian and Navon 2016) capture faithfully the DMD eigenvalues. Next, the analysis of the same flow is repeated in presence of noise. Specifically, the fluid flow is perturbed with additive white noise using a signal-to-noise (SNR) ratio of 10. Now, Figure 3b shows that the performance of the different algorithms is indeed distinct. The deterministic algorithm performs best, capturing the first 11 eigenvalues. The randomized DMD algorithm captures faithfully the first 9 eigenvalues as well as the randomized SVD based DMD algorithm does. This time, the results were obtained by computing  $q = 2$  power iterations, in addition to oversampling. The compressed DMD algorithm performs slightly worse, using the sampling parameter  $c = 500$ , however, the accuracy could be improved using a larger sampling parameter, e.g.,  $c = 5000$ .

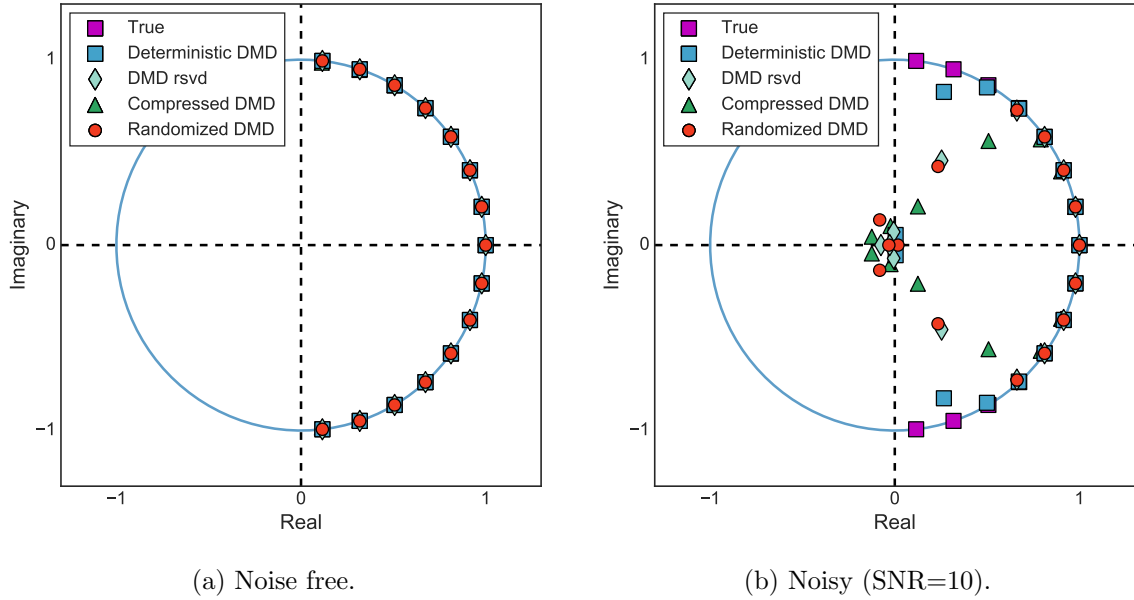


Figure 3: In absence of noise the DMD eigenvalues are faithfully captured as shown in (a). The performance of the different algorithms is distinct in presence of additive white noise as shown in (b). Here, the rDMD algorithm shows to be more robust than the cDMD algorithm.

Table 1 and Table 2 summarizes the results for both the noise free and noisy case. The proposed randomized DMD algorithm features the best trade-off between computational performance

and approximation quality. The main advantage of the randomized DMD algorithm over the cDMD algorithm is that the approximation quality can be controlled via oversampling and power iterations. The results show that the randomized DMD algorithm is more robust, and not as sensitive as the cDMD algorithm is to noise in the data. Further, the standard deviation (SD) of the reconstruction error over 100 runs indicates that the randomized algorithm is the more reliable algorithm.

Table 1: Computational performance of the deterministic and probabilistic DMD algorithms (target rank is  $k = 15$ ) in absence of noise. The results are averaged over 100 runs.

Method	Parameters	Time (s)	Speedup	Error	SD
Deterministic DMD	-	1.32	-	5.11e-03	-
DMD using randomized SVD	p=10, q=0	0.47	2.8	5.11e-03	7.55e-07
Compressed DMD	c=500	0.28	4.7	8.07e-03	1.18e-03
Randomized DMD	p=10, q=0	0.21	6.3	5.17e-03	7.02e-05

Table 2: Computational performance of the deterministic and probabilistic DMD algorithms (target rank is  $k = 15$ ) in presence of white noise. The results are averaged over 100 runs.

Method	Parameters	Time (s)	Speedup	Error	SD
Deterministic DMD	-	1.32	-	7.99e-02	-
DMD using randomized SVD	p=10, q=2	0.61	2.1	8.53e-02	1.05e-03
Compressed DMD	c=500	0.28	4.7	1.84e-01	7.57e-03
Randomized DMD	p=10, q=2	0.35	3.7	8.43e-02	1.36e-03

The computational advantage of randomized algorithms becomes pronounced with increasing dimensions. Figure 4 shows the computational times for varying target ranks, using an input matrix of dimension  $500000 \times 500$ . For clarity the cDMD algorithm was omitted here, however, randomized DMD shows a better performance for target ranks  $k < 30$ .




Figure 4: Average algorithm runtimes, over 100 runs, for varying target ranks. The input matrix is a random matrix of dimension  $500000 \times 500$ . The randomized DMD algorithms can achieve substantial computational savings.

## 6. Conclusion

Randomness as a computational strategy has recently proved capable of efficiently solving many standard problems in linear algebra. The need for highly efficient algorithms becomes increasingly important in the area of ‘big data’. We have proposed a novel randomized algorithm for computing the low-rank dynamic mode decomposition. Specifically, we have shown that the DMD can be embedded in the probabilistic framework formulated by [Halko et al. \(2011\)](#). This framework not only enables computations at scales far larger than what was previously possible, but it is also modular and flexible. Hence, it can be also utilized as a framework to embed other innovations around the dynamic mode decomposition, for instance, see [Kutz, Brunton, Brunton, and Proctor \(2016\)](#) for an overview.

The numerical results show that randomized dynamic mode decomposition (rDMD) has computational advantages over previously suggested probabilistic algorithms for computing the dominant dynamic modes and eigenvalues. The time complexity of the proposed algorithm is of order  $O(mnk)$ . [Woolfe et al. \(2008\)](#) showed that the time complexity can be reduced even further to  $O(mn \cdot \log(k))$  by using a structured random test matrix to sample the range of the input matrix, instead of a dense random test matrix. However, more importantly only two passes over the input matrix are required. This aspect is crucial in order to efficiently deal with massive data which are too big to fit into fast memory. The approximation quality of rDMD can be controlled using the concept of oversampling and the power scheme. In practice, the default values  $p = 10$  and  $q = 1$  show a good trade-off between computational performance and accuracy. The randomized DMD algorithm is, in particular, computationally beneficial if the target rank is  $k < \frac{\min\{m,n\}}{2}$ . Moreover, the presented algorithm is embarrassingly parallel, and thus the computational performance can benefit from a GPU accelerated implementation, e.g., similar to the work by [Pendergrass, Kutz, and Brunton \(2016\)](#).

Another interesting research direction is around modern randomized single pass algorithms for low-rank matrix approximation as recently proposed by [Tropp, Yurtsever, Udell, and Cevher \(2016\)](#). These algorithms can become essential, in particular, if the costs of data transfer is expensive. Further, single pass algorithms can be used to integrate DMD into a streaming model.

## References

- Bistrián D, Navon I (2016). “Randomized Dynamic Mode Decomposition for Non-Intrusive Reduced Order Modelling.” *International Journal for Numerical Methods in Engineering*, pp. 1–22. doi:10.1002/nme.5499.
- Brunton SL, Proctor JL, Tu JH, Kutz JN (2015). “Compressed sensing and dynamic mode decomposition.” *Journal of Computational Dynamics*, 2(2), 165–191. doi:10.3934/jcd.2015002.

- Colonius T, Taira K (2008). “A Fast Immersed Boundary Method using a Nullspace Approach and Multi-Domain Far-Field Boundary Conditions.” *Comp. Meth. App. Mech. Eng.*, **197**, 2131–2146.
- Drineas P, Mahoney MW (2016). “RandNLA: Randomized Numerical Linear Algebra.” *Communications of the ACM*, **59**(6), 80–90. doi:10.1145/2842602.
- Erichson NB, Brunton SL, Kutz JN (2016a). “Compressed Dynamic Mode Decomposition for Background Modeling.” *Journal of Real-Time Image Processing*, pp. 1–14. doi:10.1007/s11554-016-0655-2.
- Erichson NB, Donovan C (2016). “Randomized low-rank Dynamic Mode Decomposition for motion detection.” *Computer Vision and Image Understanding*, **146**, 40–50. ISSN 1077-3142. doi:10.1016/j.cviu.2016.02.005.
- Erichson NB, Voronin S, Brunton SL, Kutz JN (2016b). “Randomized matrix decompositions using R.” *arXiv preprint arXiv:1608.02148*.
- Gu M (2015). “Subspace Iteration Randomization and Singular Value Problems.” *SIAM Journal on Scientific Computing*, **37**(3), A1139–A1173. doi:10.1137/130938700.
- Halko N, Martinsson PG, Tropp JA (2011). “Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions.” *SIAM Review*, **53**(2), 217–288. doi:10.1137/090771806.
- Kutz JN, Brunton SL, Brunton BW, Proctor JL (2016). *Dynamic Mode Decomposition: Data-Driven Modeling of Complex Systems*. SIAM.
- Mahoney MW (2011). “Randomized algorithms for matrices and data.” *Foundations and Trends in Machine Learning*, **3**(2), 123–224. doi:10.1561/22000000035.
- Martinsson PG (2016). “Randomized Methods for Matrix Computations and Analysis of High Dimensional Data.” *arXiv preprint*, pp. 1–55. <https://arxiv.org/abs/1607.01649>.
- Noack BR, Afanasiev K, Morzynski M, Tadmor G, Thiele F (2003). “A Hierarchy of Low-Dimensional Models for the Transient and Post-Transient Cylinder Wake.” *J. Fluid Mech.*, **497**, 335–363.
- Pendergrass SD, Kutz JN, Brunton SL (2016). “Streaming GPU Singular Value and Dynamic Mode Decompositions.” *arXiv preprint arXiv:1612.07875*.
- Rokhlin V, Szlam A, Tygert M (2009). “A Randomized Algorithm for Principal Component Analysis.” *SIAM Journal on Matrix Analysis and Applications*, **31**(3), 1100–1124.
- Rowley C, Mezić I, Bagheri S, Schlatter P, Henningson D (2009). “Spectral analysis of nonlinear flows.” *Journal of Fluid Mechanics*, **641**, 115–127.
- Schmid P (2010). “Dynamic mode decomposition of numerical and experimental data.” *Journal of Fluid Mechanics*, **656**, 5–28. doi:10.1017/S0022112010001217.
- Taira K, Colonius T (2007). “The Immersed Boundary Method: A Projection Approach.” *J. Comp. Phys.*, **225**, 2118–2137.
- Tropp JA, Yurtsever A, Udell M, Cevher V (2016). “Randomized Single-View Algorithms for Low-Rank Matrix approximation.” *arXiv preprint arXiv:1609.00048*.
- Tu J, Rowley C, Luchtenberg D, Brunton S, Kutz JN (2014). “On Dynamic Mode Decomposition: Theory and Applications.” *Journal of Computational Dynamics*, **1**, 391–421.
- Woolfe F, Liberty E, Rokhlin V, Tygert M (2008). “A Fast Randomized Algorithm for the Approximation of Matrices.” *Applied and Computational Harmonic Analysis*, **25**(3), 335–366.